

Express Mail No. ER447510974US

Attorney Docket No. Pillar 715

**DATA REPLICATION
IN DATA STORAGE SYSTEMS**

Inventors: Dave Burton, Noel Otterness, & Kurt Shoens

BACKGROUND

The invention relates to data replication between primary and secondary data storage systems.

Please file 8/15/06
This application incorporates by reference herein as follows:

U.S. Application No. 10/264,603, entitled, Systems and Methods of Multiple Access Paths to Single Ported Storage Devices, filed on October 3, 2002 (Attorney Docket Pillar 701);

U.S. Application No. 10/354,797, Methods and Systems of Host Caching, filed on January 29, 2003 (Attorney Docket No. Pillar 709);

U.S. Application No. 10/397,610, Methods and Systems for Management of System Metadata, filed on March 26, 2003 (Attorney Docket No. Pillar 707);

U.S. Application No. 10/440,347, Methods and Systems of Cache Memory Management and Snapshot Operations, filed on May 16, 2003 (Attorney Docket No. Pillar 713);

U.S. Application No. 10/600,417, Systems and Methods of Data Migration in Snapshot Operations, filed on June 19, 2003 (Attorney Docket No. Pillar 711);

U.S. Application No. 10/616,128, Snapshots of File Systems in Data Storage Systems, filed on July 8, 2003 (Attorney Docket No. Pillar 714); and

U.S. Application No. 10/677,560, Systems and Methods of Multiple Access Paths to Single Ported Storage Devices, filed on October 1, 2003 (Attorney Docket No. Pillar 716).

Many enterprises require an extra copy of the data if the primary data storage system fails. Tape backup can provide the copy but is too slow for regular access to the data and is time consuming to restore to faster storage devices such as disk drives. Data replication provides a solution by transmitting a copy of the data from the storage devices of a primary data storage system to those of a secondary data storage system so that the data can be still quickly accessed when the primary data storage system fails.

High performance requires that a primary data storage system acknowledge a write command from an application before transmitting the write command to the secondary data storage system. This is referred to as asynchronous data replication. Asynchronous replication can result in the delivery of data at the secondary data storage system which is out of order with respect to the order in which the write commands were originally written by the application. When the secondary data storage system writes data out of order, the data may not match the copy at the primary data storage system.

The primary data storage system and the network (e.g. Fibre Channel) can inadvertently change the order of the data received at the secondary data storage system. A data replication system must somehow maintain a matching copy at the secondary data storage system. At the same time, data replication should keep the network busy between the primary and secondary data storage systems. Therefore, the primary data storage system should transmit write commands as soon as possible. However, many data replication systems will transmit the write commands in groups to keep the data in order resulting in inefficient use of the network.